

The ATLAS Level-1 Central Trigger

Mark Stockton^a for the ATLAS collaboration

^a*CERN,*

CH-1211 Geneva 23, Switzerland

E-mail: Mark.Stockton@cern.ch

ABSTRACT: The ATLAS Level-1 trigger system is responsible for reducing the anticipated LHC collision rate from 40 MHz to less than 100 kHz. This Level-1 selection identifies, jet, tau/hadron, electron/photon and muon candidates, with additional triggers for missing and total energy. These inputs are used by the Level-1 Central Trigger to form a Level-1 Accept decision. This decision, along with summary information, is then passed into the higher levels of the trigger system and sub-detectors, which also receive the clock from the Level-1 Central trigger. Results from commissioning the Central Trigger with cosmic rays and its performance during the first collisions will be shown.

KEYWORDS: Trigger concepts and systems (hardware and software), Trigger detectors.



Contents

1. Introduction	1
2. L1 Central Trigger	2
2.1 Hardware Overview	2
2.2 Operation	2
2.3 MUCTPI Timing	3
2.4 Monitoring	3
2.5 Bunch groups	3
2.6 Dead-Time	4
2.7 Developments	5
3. Conclusion	5

1. Introduction

The ATLAS trigger uses a three-level architecture. The level 1 (L1) trigger, is entirely implemented in custom built electronics and is designed to reduce the rate from the anticipated 40 MHz to less than 100 kHz. The subsequent higher-level trigger (HLT) selection is done by software run on PC farms. The L1 trigger decision (L1 accept or L1A) is made by the Central Trigger Processor (CTP) using information from coarse grained calorimeter information, dedicated muon-trigger detectors, and a variety of additional trigger inputs from detectors in the forward regions. A full description of the ATLAS detector can be found here [1], with a brief outline of some of the L1 inputs below.

The Liquid Argon and Tile Calorimeters send 7200 analogue signals to the L1 calorimeter trigger, which are lower granularity trigger towers, typically covering 0.1×0.1 in η, ϕ , to identify electron/photon, tau/single hadron and jet signatures along with calculating missing and total transverse energy sums. The L1 muon trigger also receives lower accuracy input from dedicated trigger detectors, chosen for fast response. Two detector types are used: Resistive Plate Chambers (RPC) in the barrel region ($|\eta| < 1.05$) and Thin Gap Chambers (TGC) in the end-caps ($1.05 < |\eta| < 2.4$). These triggers can provide up to two candidates muons per sector for six programmable thresholds, which then allows further analysis with the higher accuracy muon physics detectors (Muon Drift Tubes and Cathode Strip Chambers) in the HLT.

A different type of detector are the Beam Pick-ups (BPTX) [2] which are electrostatic button pick-up detectors located 175 m upstream of the interaction point. They are used to monitor the phase, with a precision better than 100 ps, between bunches and the LHC clock, which drives the ATLAS electronics. Originally they were used directly as beam triggers, but now they provide the LHC bunch crossing pattern used by ATLAS to create bunch groups as described in section 2.5. There are also other forward detectors, not described here, that input to the L1 trigger, these are : Minimum Bias Trigger Scintillators, Zero Degree Calorimeter, Beam Conditions Monitors and a Luminosity measurement using Cerenkov Integrating Detector.

2. L1 Central Trigger

The central trigger system forms the L1 trigger decision, following the logic defined by the trigger menu. It consists of the CTP, Muon to CTP Interface (MUCTPI) and the Trigger Timing and Control system (TTC). Here a brief overview is given of the main parts of the system (see [1] for more details) followed by details of the performance with first beams, along with new developments needed after collisions began.

2.1 Hardware Overview

Figure 1 shows diagrams of the main boards of the central trigger. The MUCTPI takes data directly from the muon detectors and forms multiplicities. It can be used to resolve overlaps and then can pass the data to the CTP. If an event passes the trigger conditions the data is also sent to the Level-2 trigger and data acquisition system (DAQ). In the CTP, the three CTP input (CTPIN) boards each accept 4×31 inputs, via LVDS signals, all of which are monitored with scalars. Of the 372 input signals 160 signals are selected by a switch matrix and made available via the Pattern In Time (PIT) bus, to then be used by the CTPCORE, see below, and for monitoring in the CTPMON, see 2.4.

The CTPCORE module forms the L1A using the signal delivered via the PIT bus as well as internally generated random triggers. It forms up to 256 trigger items by combining the trigger conditions. Some of these can be bunch groups, see section 2.5, which specify the LHC bunch numbers (time within the LHC orbit) where each particular item is allowed to fire. Individual prescales are optionally applied to each item, followed by a veto signal, then finally a logical OR of the items. The veto signal introduces dead-time, see section 2.6, using protective dead-time rules and by requests from ATLAS subsystems. The L1A, and other timing signals, are then transmitted to the sub-detector systems via the TTC. The CTP data stored upon each L1 accept (L1A) contains the status of the PIT bits and all items: before and after prescale and after veto for the firing bunch, and optionally in a window of early and late signals.




Figure 1. Outline of the a) MUCTPI and b) CTP boards.

2.2 Operation

The central trigger system is installed underground at ATLAS and fully operational. It runs with L1 output rates of typically around 10^4 Hz, although it is capable of handling much higher rates. Current ATLAS runs use a trigger menu of 230 out of the 256 items available and the prescale sets being updated on average seven times per run, to adjust the trigger rates to the fixed output bandwidth with decreasing luminosity. The CTP is maintained by a group of roughly ten people, with one on call at all times usually receiving less than one call per week on average.

2.3 MUCTPI Timing

One of the checks performed on the MUCTPI was to check for transmission errors by performing a clock fine delay scan between the MUCTPI and the logic modules of the RPC and TGC. The test indirectly measures the relative phase between the incoming muon trigger sector data and the MUCTPI clock. It is performed by altering the phase of the MUCTPI clock, that strobes the incoming muon sector data, by 0.5 ns steps over the full 25 ns range. During this the sector logic modules send a known repetitive test pattern and then for each delay step the data transmission is checked using diagnostics memories. Figure 2 shows that with the current operating point at 3 ns (for the MUCTPI clock fine delay), the signals are strobed correctly with no errors and with timing margins of more than ± 5 ns for all 208 trigger sectors.

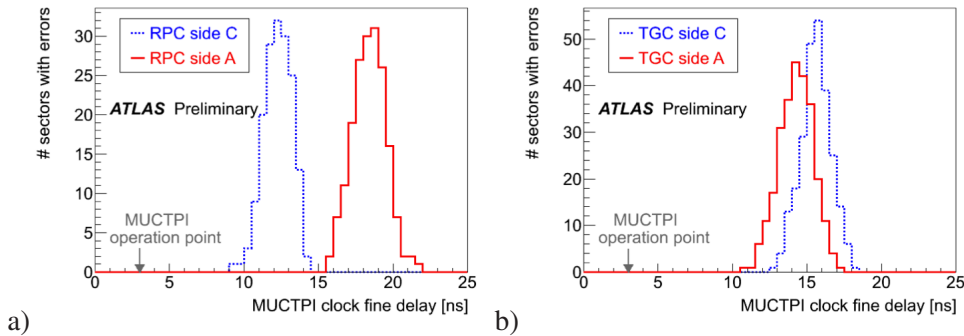


Figure 2. Results from a clock fine delay scan, showing the number of sectors with at least one error per delay setting, for a) RPC and b) TGC.

2.4 Monitoring

The rates of the 256 trigger items are monitored with scalers before and after prescale, and after veto. The CTPMON is equipped with scalers counting the number of signals for all 3564 bunches (each representing a 25 ns portion of the beam orbit), for each of the 160 PIT lines. An example of this per-bunch monitoring is shown in section 2.6. The counters are read out and published as online monitoring histograms every ten seconds. The numbers are also written to a conditions database for permanent storage at a configurable rate, typically once per five minutes. The CTPIN is equipped with 768 counters to monitor all inputs, but unlike the CTPMON they are not bunch-aware. The CTPIN data-rate is thus much lower than that of CTPMON and can be published and permanently stored with a higher frequency. In addition, all inputs can be monitored, not only those available at the PIT bus.

2.5 Bunch groups

The CTP bunch group conditions are used in combinational logic “AND” with other trigger conditions to define which items generate a L1A. There are eight distinct groups, each with their own particular purpose, defined for each LHC bunch. An illustrative example is given in figure 3. The *collision* bunch group contains the bunch numbers for which the two beams meet in ATLAS, and

is thus used by physics data trigger items. The three *single beam* bunch groups, used for beam background estimates, select bunches corresponding to beams passing through ATLAS, but not colliding. Another bunch group selects a few *bunches after collisions*, again for monitoring backgrounds and also for slow particle searches. Two bunch groups have a more technical purpose; the *calibration requests* group defines the times at which sub-detectors may request calibration triggers, typically in the long gap with no collisions. The *bunch count reset veto* leaves a short time slice for distribution of the LHC bunch count signal to the on detector electronics. As the LHC fill scheme can vary from fill to fill, ATLAS has developed and commissioned a procedure for monitoring and redefining the bunch groups using the BPTX, see section 1. An online application measures the fill scheme seen by the BPTX and calculates the corresponding bunch groups. The ATLAS trigger shifter compares the suggested bunch groups to the current configuration of the CTP and may generate a new configuration at the press of a button. Changes to the central trigger configuration are protected by a short veto and stored to conditions database.

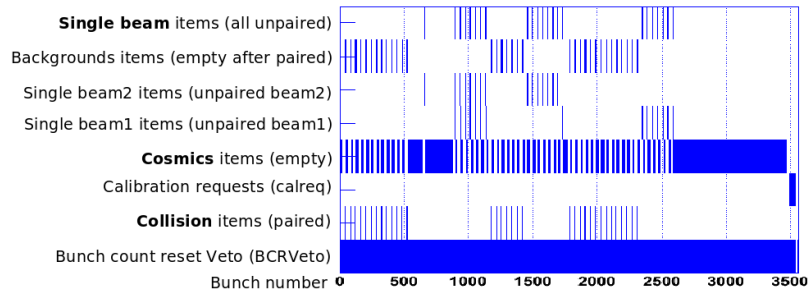


Figure 3. Example bunch groups for 50 bunches, with 30 colliding in ATLAS. The filled area shows where the appropriate group is applied, see the text for more details of the groups.

2.6 Dead-Time

Preventive dead-time is introduced by the CTP to stop the front-end buffers from overflowing, and is calculated in two ways: simple (or fixed) and complex. The simple dead-time is a programmable number of bunch crossings after each L1A, used for example to avoid overlapping readout windows. The complex dead-time uses a leaky bucket model to emulate a front end buffer. In this model when the bucket is full there is dead-time. It is defined by X (in units of L1A) being the size of the bucket (i.e. the front-end buffer) and R (in BC) being the time it takes to leak 1 L1A. With these numbers the trigger rate, on average, is limited to X triggers in a time period of $X \times R$ bunch crossings. The current settings within ATLAS are: 5 BC (simple) 7 in 415BC (complex). As already described in section 2.4, the dead-time is monitored per bunch in the CTPCORE as can be seen in figure 4. This shows the bunches that produce collisions in ATLAS from an LHC filling scheme. This then produces a rate of L1A, which is matched by the simple dead-time veto. The complex dead-time has a more complicated curve, showing how the front end buffers fill and empty. There is also the DAQ busy coming from the sub-detector back-ends that is in logical “OR” with the preventive dead-time.

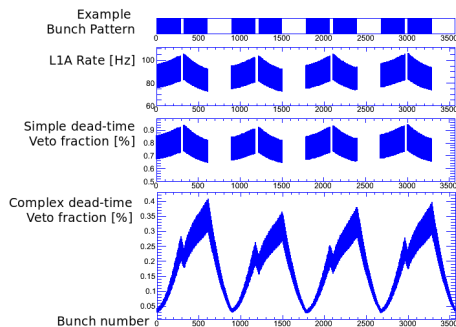


Figure 4. Example of a close to nominal LHC bunch pattern, LIA rate and dead-time (both simple and complex) veto, as measured by the CTP counters.

2.7 Developments

With the arrival of collisions came ideas for new developments of which four examples are explained below. A zero bias trigger is implemented in the hardware of the CTPIN, to be used for comparing the normalization of collisions and simulations. It uses a specific input trigger to fire the trigger exactly one LHC turn after this input trigger fired. The input trigger is chosen to deliver a completely unbiased bunch-aware trigger signal with a rate proportional to the luminosity. Currently it uses a 10 GeV electromagnetic trigger. Also now implemented in the hardware are fractional prescales, in the CTPCORE. The advantage of fractional prescales, over integer prescales, is that rates can be scaled more appropriately, i.e. an item with prescale 1 can be scaled by 10% rather than just changing the prescale to 2. Expanding on this further is the automatic rate prescaling where the CTPCORE software checks certain monitoring (i.e. not physics) trigger items against configurable threshold values. Should the rate exceed the threshold then the prescale is updated, after approval from the shifter. Should the rate lower again then the prescale will automatically revert, again after approval from the shifter. Another automatic feature added to the software was for enabling inputs. This is part of the CTPIN software and automatically masks inputs from detectors that are disabled, which, for example, can be the case during ATLAS runs with parallel sub-detector tests or calibration runs.

3. Conclusion

The ATLAS L1 trigger has been fully commissioned. The BPTX have been used successfully to identify the filling pattern used to define the bunch groups. New CTP features have been developed, along with improved monitoring, which are now in use. The CTP has performed flawlessly and reliably, to allow ATLAS to collect a 7 TeV physics data-set of 16.72 pb^{-1} (on 13/10/2010.).

References

- [1] The ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, 2008 JINST **3** S08003.
- [2] C. Ohm, T. Pauly, *The ATLAS beam pick-up based timing system*, *Nucl. Instr. and Meth. A*(2010), arXiv:0905.3648v1